

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
22 November 2001 (22.11.2001)

PCT

(10) International Publication Number  
**WO 01/88834 A2**

(31) International Patent Classification: G06K 9/00

(21) International Application Number: PCT/US01/15140

(22) International Filing Date: 10 May 2001 (10.05.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/204,816 17 May 2000 (17.05.2000) US

(71) Applicant (for all designated States except US): NEW YORK UNIVERSITY [US/US]; 70 Washington Square South, New York, NY 10012-1091 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): BERGER, Gideon

[US/US]: Apartment 4G, 50 East 89th Street, New York, NY 10128 (US). MISHRA, Bhubaneswar [IN/US]; 16 Dunster Road, Great Neck, NY 11021 (US).

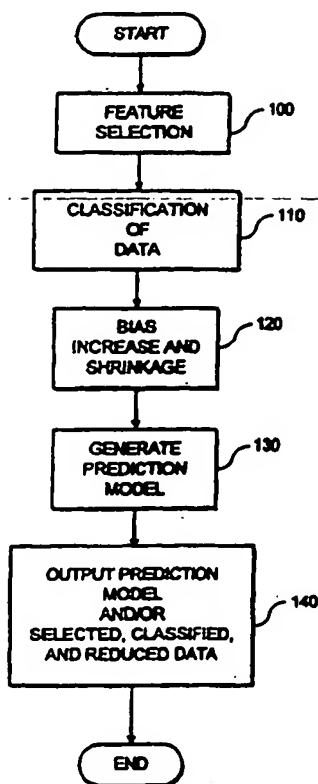
(74) Agents: CLARK, Richard, S. et al.; Baker Botts LLP, 30 Rockefeller Plaza, New York, NY 10112-0228 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR DATA CLASSIFICATION IN THE PRESENCE OF A TEMPORAL NON-STATION-ARITY



(57) Abstract: A method and system for determining a feature of a particular pattern are provided. In particular, data records are received, and predetermined patterns that are associated with at least some of the data records are obtained. Using the system and method, particular information is extracted from at least a subset of the received data records, the particular information being indicative of the particular pattern in at least some of the data records. Then, it is determined whether the particular pattern is an unexpected pattern based on the obtained predetermined patterns. In addition, it is possible to classify and reduce data and/or parameters provided in the data records. First, the data records are received. Then, the data records which have at least one particular pattern are classified using a Multivariate Adaptive Regression Splines technique. Thereafter, the data and/or parameters of the classified data records are shrunk using a Stein's Estimator Rule technique.



IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**Published:**

— without international search report and to be republished upon receipt of that report

## METHOD AND SYSTEM FOR DATA CLASSIFICATION IN THE PRESENCE OF A TEMPORAL NON-STATIONARITY

### Field of the Invention

The present invention relates to a method and system for classifying data, and  
5 more particularly to a data classification method and system in the presence of a  
temporal non-stationarity.

### Background Information

Approaches for predicting the value of a dependent response variable based  
the values of a set of independent predictor variables have been developed by  
10 practitioners in the art of the statistical analysis and data mining for a number of  
years. Also, a number of conventional approaches for modeling data have been  
developed. These known techniques require a set of restrictive assumptions about the  
data being modeled. These assumptions include, e.g., a lack of noise, statistical  
independence, time invariance, etc. Therefore, if the real data being modeled is  
15 dependant on certain factors which are contrary to the assumptions required for the  
accurate modeling by the conventional techniques, the results of the above-described  
conventional data modeling would not be accurate.

This is especially the case in the presence of temporal, non-stationary data.  
Indeed, no robust approach which considers such data has been widely used or  
20 accepted by those in the art of the statistical analysis. For a better understanding of  
the difficulties with the prior art approaches, temporal data and non-stationary data  
are described below.

Temporal data refers to data in which there exists a temporal relationship  
among data records which varies over time. This temporal relationship is relevant to  
25 the prediction of a dependent response variable. For example, the temporal data can  
be used to predict the future value of the equity prices, which would be based on the  
current and past values of a set of particular financial indicators. Indeed, if one  
believes in the importance of trends in the market, it is not enough to simply consider

the current levels of these financial indicators, but also their relationships to the past levels.

In another example, the supermarket application may prefer to group certain items together based on the purchasers' buying patterns. In such scenarios, the temporal data currently used in such supermarket application is the data provided for each customer at the particular checkout, i.e., a single event. However, using the data at the checkout counter for a single customer does not take into consideration the past data for this customer (i.e., his or her previous purchases at the counter). In an example of an intrusion detection system, the use of the time-varying data is very important. For example, if a current login fails because the password was entered incorrectly, this system would not raise any flags to indicate that an unauthorized access into the system is being attempted. However, if the system continuously monitors the previous login attempts for each user, it can determine whether a predetermined number of failed logins occurred for the user, or if a particular sequence of events occurred. This event may signify that an unauthorized access to the system is being attempted.

Non-stationary data refers to data in which the functional relationship between the predictor and response variables changes when moving from in-sample training data to out-of-sample test data either because of inherent changes in this relationship over time, or because of some external impact. For example, with a conventional network intrusion detection system, a predictive model of malicious network activity can be constructed based on, e.g., TCP/IP log files created on a particular network, such as the pattern formed from the previous intrusion attempts. However, intruders become more sophisticated in their attack scenarios, attack signatures will evolve. In addition, the conventional intrusion detection systems may not be usable for all conceivable current operating systems, much less for any future operating systems. An effective intrusion detection system must be able to take into consideration with these changes.

One of the main difficulties being faced by the conventional predicting engines is that the data is "multi-dimensional" which may lead to "over-fitting". While it is possible to train the prediction system to make the predictions based on the

previous data, it would be difficult for this system to make a prediction based on both new data and the data which was previously utilized to train the system. The conventional systems utilize predictor values for each category of the data so as to train themselves as described above. For example, if the prediction system intends to predict the performance of certain baseball teams, it would not only use the batting average of each player of the respective team, but also other variables such as hitting powers of the respective players, statistics of the team while playing at home, statistics of the team when it is playing away from home, injury statistics, age of the players, etc. Each of these variables has a prediction variable associated therewith. Using these prediction variables, it may be possible to train the system to predict the performance of a given baseball team.

However, the conventional systems and methods described above are not flexible enough to perform its predictions based on a new variable (e.g., the number of player leaving the team) and a new corresponding prediction variable being utilized for the analysis. In addition, it is highly unlikely that the data values being utilized by the conventional systems and methods, i.e., after the system has already been trained, is the same as or similar to the data of the respective prediction variables that were already stored during the training of this system. The above-described example illustrates what is known to those having ordinary skill in the art as "over-fitting". As an example to illustrate this concept, the system may only be trained using training data (e.g., in-sample data) which can represent only 0.1% of the entire data that this system may be required to evaluate. Thereafter, the prediction model is built using this training data. However, when the system is subjected to the real or test data (e.g., out of sample data), there may be no correlation between the training data and the real or test data. This is because the system was only subjected to training using a small portion of the real/test data (e.g., 0.1%), and thus never seen most of the real or test data before.

There is a need to overcome the above-described deficiencies of the prior art systems, method and techniques. In particular, there is a need to provide a method and system for classifying data that is temporal and non-stationary.

**Summary of the Invention**

A classification system and method according to the present invention offers an approach for a prediction in the presence of temporal, non-stationary data which is advantageous over the conventional systems and methods. The first exemplary step of the system and method uses temporal logic for discovering features provided in the data records. The next exemplary step is the classification of the data records using the selected features. Another exemplary step of the system and method of the present invention utilizes a "shrinkage technique" to reduce the undesirable effect of "over-fitting".

Accordingly, a method and system according to the present invention are provided for determining a feature of a particular pattern. Using these exemplary system and method, data records are received, and predetermined patterns that are associated with at least some of the data records are obtained. Using the system and method, particular information is extracted from at least a subset of the received data records, the particular information being indicative of the particular pattern for at least some of the data records. Then, it is determined whether the particular pattern is an unexpected pattern based on the obtained predetermined patterns. At least one record of the data records may include temporal data and/or non-stationary data.

In another embodiment of the present invention, the predetermined patterns are obtained by assigning a threshold, and correlating the data records into sets of patterns as a function of the threshold. Also, the determination of whether the particular pattern is an unexpected pattern include a determination if the particular pattern corresponds to at least one pattern of the sets of patterns. The positive determination regarding the unexpected pattern can be made if the particular pattern does not correspond to any pattern of the sets of patterns.

In yet another embodiment of the present invention, the unexpected pattern can be indicative of an interestingness measure in the predetermined pattern. In addition, the data records can include input sequences, and the input sequences can be scanned to determine an interestingness measure of at least one event in the input sequences. It is also possible to initialize a pattern list by inserting all events of the input sequences therein. Then, from all patterns in the pattern list, a first pattern which

has a largest interestingness measure may be selected. The data records may include a maximum allowable length value. Thus, the first pattern can be expanded to be a second pattern. If a length of the second pattern is greater than the maximum allowable value, the second pattern can be added to the pattern list. Thereafter, if a  
5 length of the second pattern is less than or equal to the maximum allowable value, the first pattern can be subtracted from the pattern list. These steps can be repeated until the pattern list becomes empty. Finally the particular pattern which includes the interestingness measure can be output.

According to still another embodiment of the present invention, a pattern list  
10 may be initialized by inserting all events of the input sequences therein, and at least one suffix list can also be initialized. Locations of certain patterns of the input sequences can be calculated, and previously discovered may be updated patterns based on the calculated locations. The pattern list of the certain patterns can then be updated. The data records can include a maximum allowable length value.

15 In another embodiment of the present invention, further records are generated by modifying the data records to include additional features. Also, a functional model is generated using the further records. A plurality of sets of the further records are also generated, and the prediction model is generated for each set of the further records. Furthermore, a single model can be generated based on each functional model of the  
20 respective set of the further records.

According to yet another embodiment of the present invention, the data  
records which have the unexpected pattern can be classified. Thereafter, a prediction model is generated as a function of the classified data records. The classification of the data records can be performed using a Multivariate Adaptive Regression Splines  
25 technique. Then, data and/or parameters of at least one of the classified data records is shrunk so as to determine a mean of the data and/or the parameters. The shrinking technique can be a Stein's Estimator Rule technique.

#### **Brief Description of the Drawings**

For a more complete understanding of the present invention and its  
30 advantages, reference is now made to the following description, taken in conjunction with the accompanying drawings, in which:

Figure 1 is an exemplary embodiment of a classification system according the present invention;

Figure 2 is a top level diagram of an exemplary embodiment of a method according to the present invention, which can be performed by the classification  
5 system of Figure 1;

Figure 3 is a flow diagram of a first exemplary feature selection technique of the method according to the present invention which performs the feature selection by utilizing a threshold to determine whether a particular pattern is an unexpected  
pattern;

10 Figure 4A is a flow diagram of a second exemplary feature selection technique of the method according to the present invention which performs the feature selection based on an interestingness measure;

Figure 4B is a flow diagram of a third exemplary feature selection technique of the method according to the present invention which performs the feature selection  
15 based on suffix lists;

Figure 5 is an illustration of an exemplary implementation of the system and method of the present invention by an intrusion detection system;

Figure 6 is a flow diagram of the exemplary embodiment of the method of the present invention utilized by the intrusion detection system of Figure 5, in which a  
20 prediction model is generated;

Figure 7 is another flow diagram of the exemplary implementation of the method of the present invention by the intrusion detection system of Figure 5;

Figure 8 an illustration of an exemplary implementation of the system and method of the present invention by a disease classification system; and

25 Figure 9 is a flow diagram of the exemplary implementation of the method of the present invention by the disease classification system of Figure 8.

### Detailed Description

Figure 1 illustrates an exemplary embodiment of a classification system 10 according to the present invention. In this drawing, the system 10 is connected to one  
30 or more databases 20 for receiving an ordered set of data records. Each data record preferably includes a set of features that may be relevant (given particular domain



knowledge) for predicting the value of a defined dependant variable. In addition, a particular data record may also include certain relationships between itself and other data records.

Upon the receipt of these data records, the system 10 according to the present invention selects and/or extract certain features from the data records, as shown in  
5 step 100 of Figure 2, which illustrates an exemplary embodiment of the method according to the present invention. These features may be temporal features that are most relevant for predicting the value of the dependent variable. Then, in step 110, the system 10 uses the method of the present invention to classify and modify the data  
10 records received from the databases 20 based on the features that were extracted from the data records and the classification thereof. Since the classified data records being generated by step 110 are numerous, it is beneficial to shrink them. (step 120 of Figure 2). Thereafter, the data records that were selected as including or being part of the particular patterns (when classified and shrunk) are used to generate a predictive  
15 model in step 130 of Figure 2. Finally, the prediction model and/or the shrunk data records and patterns are output. For example, Figure 1 illustrates that such output can be provided to a printer 30 for generating hard copies of the predicted model or shrunk data, forwarded to a display device 40, stored on a storage device 50, and/or transmitted via a communications network 60 to another device (not shown in  
20 Figure 1).

According to one exemplary embodiment of the present invention, the system 10 can be a general purpose computer (e.g., Intel processor-based computer), a special purpose computer, a plurality of each and/or their combination. The storage device 50 can be one or more databases, one or more hard drives (e.g., stackable hard drives)  
25 internal RAM, etc. The communications network 60 can be the Internet, Intranet, extranet, or another internal or external network. It is also within the scope of the present invention to receive the data records from the databases 20 via a communications network, such as the Internet, Intranet, etc. The details of exemplary embodiments of the present invention are provided below.

## I. FEATURE SELECTION

To accomplish the extraction/selection of the features from the data records, the classification system 10 searches and preferably selects certain patterns in the data records which can be defined as having an "interestingness measure". This particular  
5 interestingness measure used is preferably domain dependent, and in general, it is the measure of how much the occurrence of the pattern correlates with the occurrence of a single value of the predicted variable. The determination of the interestingness measure can be useful in a number of examples, such as, e.g., for a network intrusion detection. When searching for patterns that characterize malicious activity on the  
10 network, not only the patterns that occur frequently in the presence of an attack are monitored, but also the selection of those patterns which occur more frequently during an attack than during the normal network activity.

The above-described example defines at least one "interestingness feature" which can be used by the system and method of the present invention for monitoring  
15 the patterns of the data records having this measure, and selecting the corresponding patterns therefrom. For example, the interestingness measure for the network intrusion system may be a ratio of a number of occurrences of the particular pattern during the course of intrusion to the number of occurrences of this pattern during the course of normal network behavior. This interestingness measure, unlike the  
20 frequency, enables an identification of patterns that are non-frequent and yet highly correlated with intrusive behavior, and provides a way to ignore patterns which occur frequently during an intrusion, but occur just as frequently during normal behavior.

In another example of the network intrusion detection, the dependent variable that may be used for the interestingness feature can have a value between 0 and 1,  
25 which represents the probability that the associated data record that can be a part of the intrusion. In this exemplary case, the interestingness measure of a pattern P is denoted as:

$$I(P) = \Pr(\text{Intrusion} \mid P).$$

The interestingness measure of the pattern P would, in this case, be the  
30 probability that the particular data record is part of the intrusion given that the pattern

P occurred. Using a predefined interestingness threshold T, the following sets of patterns can be included in the data records as additional features:

$$S1 = \{P | I(P) > T\}, S2 = \{P | I(P) < 1-T\}, S3 = \{P | \neg P \in S2\}$$

For example, set S1 may represent the most interesting patterns. In the case of the intrusion detection, set S1 may be defined as a set of patterns that are most highly correlated with the intrusion based on the training of the prediction model using in-sample data. Set S2 may include the least interesting patterns, or in the exemplary intrusion detection, set S2 may represent the most highly correlated patterns with a normal behavior also based on the training of the prediction model using in-sample data. Set S3 may have the patterns whose negation is provided in set S2. The purpose of set S3 is to aid in the mitigation of the effects of non-stationarity.

For example, in the intrusion detection scenario, the system 10 and method according to the present invention take into consideration the situation in which the out-of-sample data set contains an intrusion that was not present in the in-sample data on which the model was based. Thus, as illustrated in Figure 3, an exemplary embodiment of the present invention provides that the system 10 receives an ordered set of data records which includes the data records used for accessing the network (step 200), and assigns a predetermined interestingness threshold T to be applied to these data records (step 210). The data records are then correlated so that particular sets of patterns are associated therewith, based on the threshold T (step 220). In step 230, it is then determined whether the current pattern (e.g., a predetermined number of unsuccessful logins to the network) corresponds to the first type of an expected event that is provided in set S1. It would not be expected that the patterns that are part of this novel attack to be in set S1, since set S1 contains the patterns associated with only those attacks present in the training data (e.g., which used the in-sample data for generating the prediction model). If the current pattern corresponds to the patterns in set S1, then the pattern is assigned as being of the first type in step 240, i.e., definitely an intrusion attack on the network. Otherwise, it is determined (in step 250) whether the current pattern corresponds to the second type of an expected event that is provided in set S2.

If the current pattern corresponds to the patterns in set S2, then the pattern is assigned as being of the second type in step 260, i.e., definitely not an intrusion attack on the network. It would not be expected that the patterns that are part of this novel attack to be in set S2 because this set S2 contains the patterns that are associated with a normal behavior of the network (as trained by the in-sample data). However, if the current pattern does not correspond to set S1 or set S2, then there is a pattern that does not neatly fit into any known set of patterns, i.e., thus being a novel attack. This pattern would not be considered as being a normal behavior on the network. According to this exemplary embodiment of the system and method according to the present invention, the pattern(s) present in the above described novel attack are considered as deviating from the patterns provided in set S2. Therefore, the current pattern has to be the third type of event, i.e., an unexpected (or interesting) event, which should be part of the set S3 of patterns that were in neither set S1 nor in set S2. Thus, in step 270, the current pattern is set as including an interestingness feature so as to identify its behavior as deviating from what is considered as the normal behavior on the network, even if this deviant behavior is not part of any known attack. After the current pattern is set as described above with reference to steps 240, 260, 270, the determination regarding the type of the event (of the current pattern) is output in step 280.

Given that the data records are populated with both a set of basic features as well as the derived features, namely temporal patterns, a classifier based on this data can be generated.

From the above described exemplary method of the present invention, it should be understood that an interestingness measure for the patterns of the data records could be defined as marking such patterns "unexpected" patterns. To find unexpected patterns, it may be preferable to first define these patterns in terms of temporal logic expressions, in sequences of the data records. For example, it is possible to assume that each event in each data record in the sequence occurs with some probability, and that certain conditional distributions on the neighboring events are present. Based on such predicates, it is possible to compute an expected number of occurrences of a certain pattern in a sequence. If the actual number of the

-11-

occurrences of a particular pattern significantly differs from the expected number of the occurrences, then this particular pattern would be considered "unexpected" and therefore interesting.

To determine the expected number of the occurrences of the particular pattern P, it may be preferable to assign a probability distribution over the events according to one exemplary embodiment of the present invention. In general, certain problem domains may suggest a preferable technique to evaluate these expectations rather than by calculating them as a function of the frequencies of individual events. In the exemplary network intrusion detection setting, it is possible to calculate the expected number of the occurrences of the particular pattern P during the attack on the network based on the frequency of the particular pattern P during the normal activity on the network. In other settings, i.e., different than the network intrusion detection, other techniques for determining the expectations may be appropriate. The underlying issue solved by the system and method of the present invention is whether given any technique for computing the expectations for the particular pattern, it is possible to efficiently identify interesting or unexpected patterns using the retrieved data records.

In one exemplary technique of the method according to the present invention, all unexpected patterns can be found if, e.g., the ratio of the actual number of occurrences to the expected number of occurrences exceeds a certain threshold. This exemplary technique is illustrated in Figure 4A. First, input string(s)/sequence(s) 305, event probabilities 306, a threshold T for the interestingness measure 307 and a number for a maximum allowable pattern length ("MAXL") 308 are provided to the system 10. The event probabilities 306 may be determined for each atomic event. The threshold T 307 may be a value that, if exceeded by the interestingness measure of a pattern, deems the pattern to be interesting. It is also possible to input a user-defined constant to the system 10 which determines the maximum number of events that a particular event or data record can precede another event or data record. Then, in step 310, the input string(s)/sequence(s) are scanned to determine the interestingness measure of each event therein. In step 315, a list L that includes all these events is initialized. From all patterns provided in the list L, a particular pattern

C is selected which has the largest interestingness measure to be the next pattern for expansion (step 320).

Then, in step 325, this particular pattern C is indeed expanded by scanning the input string(s)/sequence(s) to detect the occurrences of the particular pattern C. When  
5 the occurrence of the pattern C is detected, the particular pattern C is expanded as a prefix and as a suffix, i.e., record all occurrences of: (C Op X) and (X Op C), where X is also a pattern, "Op" ranges over the temporal operators, and X ranges over all events. Thereafter, the interestingness or unexpected pattern(s) of all newly discovered patterns C' is determined, i.e., by the system 10 as described below.

10 In step 330, it is determined whether the length of the newly discovered patterns C' is smaller than the maximum allowable length (MAXL, and if so, the newly discovered patterns C' can be removed from the list L (step 340). Otherwise, the particular pattern C is removed from the list L in step 335. In step 345, it is determined whether the list L is empty. If not, the processing of this exemplary  
15 technique of the method according to the present invention is returned to step 320. Otherwise, in step 350, the interesting pattern(s) are output by the system 10, e.g., to the printer 30, the display device 40, the storage device 50 and/or the communications network 60.

In another exemplary embodiment of the present invention, it is possible to  
20 start with small patterns, and expand only those patterns that offer the potential of leading to the discovery interesting/unexpected, larger patterns. Using this exemplary technique, it is preferable to first find all patterns that occur relatively frequently, given a class of operators, an input sequence of events, and a frequency threshold. The exemplary technique for solving this problem has two alternating phases:  
25 building new candidate patterns, and counting the number of occurrences of these candidates.

The efficiency of this exemplary technique is based on two observations:

- a. Where there are potentially a large number of patterns that have to be evaluated, the search space can be dramatically pruned by building  
30 large patterns from smaller ones in a prescribed way. For example if a pattern " $\alpha N \beta N \gamma$ " is frequent, then the patterns " $\alpha N \beta$ " and " $\beta N \gamma$ "

must also be frequent. Thus, for a pattern P to be frequent, its sub-patterns should also be frequent. The exemplary technique for identifying frequent patterns can take advantage of this fact by considering the patterns of size n if its prefix and suffix of size n-1 are themselves frequent.

- b. All complex patterns can be the result of recursively combining other smaller patterns. For example, in order to efficiently count the number of occurrences of the pattern " $\alpha N \beta B_K \delta B_K \gamma$ ", it is preferable to identify the number of occurrences and location of the two patterns " $\alpha N \beta$ " and " $\delta B_K \gamma$ ", and to have an efficient way for combining the patterns via the  $B_K$  operator. In general, since all of exemplary operators can be binary, when combining two patterns with operator Op to create a larger pattern and determine the number of occurrences of the resulting pattern, it is preferable to determine the number and locations of Op's two operands, and to provide an efficient way for locating patterns of the form A Op B.

The exemplary technique according to the present invention initially counts the number of occurrences of length 1 patterns (e.g., the length of the pattern is the number events that occur in it). Thereafter, a candidate set for the next iteration of discovery is computed by combining, in a pair-wise manner all frequent length-1 patterns via each operator. For example, in the nth iteration, the combination of the patterns of length n-1 and length 1 can be added to the candidate set provided that the length (n-1) prefix and suffix of the resulting length n pattern have already been deemed frequent in the previous iteration. Then, during the discovery phase, the number and location of the occurrences of the candidate length n patterns can be determined given the locations of their length n-1 prefixes and length 1 suffixes. This process continues until the candidate set (or list) becomes empty. The memory requirements of this exemplary technique are minimized because once a pattern is deemed as being infrequent, it can never result in being the sub-pattern of a larger frequent pattern, and can therefore be discarded. Such property may not hold in view of the definition of interestingness provided above, as shall be discussed in further

detail below. In particular, a pattern can be unexpected while its component sub-patterns may be expected. This feature of the interestingness measure can be understood using the following example:

Let the set of events be  $E = \{A, B, C\}$ . Assume that the probability of these events is  $\Pr[A]=0.25$ ;  $\Pr[B]=0.25$ ; and  $\Pr[C]=0.50$ . Also assume that these events are independent. Let the interestingness threshold  $T=2$ , i.e., for a pattern to be interesting, the value of the actual number of occurrences of the pattern divided by the expected number of occurrences of the pattern should preferably exceed 2. For example, the following string of events can be input into the system 10:

10 ABABABABCCCCCCCCCCCC (the length of this string being  $N=20$ )

Given the above-mentioned probabilities,  $E[A]=5$  and  $E[B]=5$ , and the expression for computing expectations for patterns of the form ANB.

$$\begin{aligned} E[ANB] &= \Pr[A]\Pr[B](N-1) \\ &= (0.25)(0.25)(19) \\ 15 &= 1.1875 \end{aligned}$$

Since  $A[A]=4$  and  $A[B]=4$ , both of the events A and B are not interesting (in fact, the actual number occurrences of these events was less than what was expected), but the pattern ANB which occurred 4 times was interesting with

$$\begin{aligned} IM(ANB) &= \frac{4}{1.1875} \\ 20 &= 3.37 \end{aligned}$$

This lack of monotonicity in the interestingness measure can result in a significantly more complex problem, specifically in terms of space complexity. In the exemplary technique for discovering frequent patterns, significant pruning of the search space may occur with each iteration. That is, when a newly discovered pattern is found to have occurred fewer times than the frequency threshold, it may be discarded as adding new events to it, and thus cannot result in a frequent pattern (which is not the case using the interestingness measure). The addition of an event to an uninteresting pattern can result in the discovery of an interesting pattern being created. This inability to prune the discovered patterns leads to a large increase in the amount of space required to find unexpected patterns.



Another exemplary technique of the method according to the present invention for finding unexpected patterns involves sequential scans over the string of events discovering new patterns with each scan is illustrated in Figure 4B. To summarize this exemplary technique, a list is maintained of those patterns that were discovered previously, and on each subsequent iteration of this technique, the "best" pattern is selected from this list for expansion to be the seed for the next scan. Described below is an exemplary method to determine which pattern is the "best" pattern.

The "best" pattern can be defined as a pattern that is most likely to produce an interesting pattern during the expansion. By expanding the already interesting pattern, it is possible, and even likely, to discover additional interesting pattern(s). However, it should still be determined which is the best candidate for the expansion among interesting patterns already discovered. If no interesting patterns remain unexpanded, it is determined whether there any uninteresting patterns worth expanding.

According to this exemplary embodiment of the present invention, input string(s)/sequence(s) 355, event probabilities 356, a threshold T for the interestingness measure 357, a number for a maximum allowable pattern length ("MAXL") 358 and a value "MIN\_TO\_EXPAND" 359 are provided to the system 10. The MIN\_TO\_EXPAND value is preferably the minimum threshold of expected interestingness that the pattern should have in order to become the next pattern. Then, a scan of the input string(s)/sequence(s) takes place, in which the number of occurrences (and therefore, the frequencies) of individual events are counted to determine the interestingness and location of each event (step 360). This scan (e.g., a linear scan) is a scan of the "DL" events that occur in the record string(s)/sequence(s), where "D" is the number of data records and "L" is the number of fields in each data record.

In step 365, the list of patterns is initialized with the set of discovered patterns. For example, certain R lists should be initialized at this stage, where R is the number of temporal operators that are used. Each list may represent the pattern form X, where X is an arbitrary literal. One sorted list can be stored for each temporal operator. The processing time and capacity preferable for this initialization

corresponds to the processing time and capacity of sorting these lists. Initially, all lists can be sorted in an identical order. Therefore, the total processing time and capacity of this initialization may be defined by  $O(N \log N)$ , where  $N$  is the number of distinct events in the database. Each literal  $\alpha$ , in each list, has an initial candidacy value of:

$$\frac{A[\alpha]}{P[\alpha]}$$

where  $A[\alpha]$  is the number of occurrences of  $\alpha$  which can be determined in the initial scan.

Then, in step 370, the suffix lists are initialized. For example, the "R" lists are preferably initialized at this stage, where  $R$  is the number of temporal operators that can be predefined or defined by a user. Each such list contains the potential suffixes for all length 2 patterns. Each of these lists would again be sorted based on their candidacy values. Initially, these candidacy values are the same as those for the set of discovered patterns (described above for step 465), and therefore no additional sorting is necessary. The total processing time and capacity of this initialization can be defined as  $O(N)$ .

In step 375, the pattern locations are calculated. As described above, it is possible to compute the locations of the pattern resulting from combining the pattern  $P$  with a literal  $\alpha$  via the operator "Op" via the linear scan of the location lists for the pattern  $P$  and the literal  $\alpha$ . The total number of operations that should be performed for this computation is proportional to the longer of these two location lists. This has an expected value of:

$$\frac{DR}{N},$$

where  $D$  is the number of data records,  $R$  is the number of temporal operators, and  $N$  is the number of distinct events in the database.

Then, the already discovered patterns are updated in step 380. Given that the locations of the candidate  $P \text{ Op } \alpha$  have been previously computed, this step entails two substeps. In the first substep, the newly discovered patterns are inserted into the

the printer 30, the display device 40, the storage device 50 and/or the communications network 60, and the processing of this exemplary method is completed.

For example, the exemplary technique described above with reference to Figure 4B continues to expand best candidates of the unexpected patterns until there are no more candidates that are worthy of expansion. To further explain this concept, the following definitions can be utilized:

Definition I: The FORM(P) of the pattern P is a logical expression with all ground terms in the pattern P replaced by variables. For example, if  $P = \alpha N \beta B_K \gamma B_K \delta$ , then  $FORM(P) = WNXB_KYB_KZ$ .

Given the length of the input string(s)/sequence(s), it is possible to determine the number of patterns of each form in the input string/sequence. For example, given a string of length M, the number of patterns of form XNY is M-1. The number of patterns  $XB_KY$  is  $(M-K)K + ((K)(K-1)/(2))$ .

Definition II: Given the pattern P and an operator Op, Actual Remaining (P Op X) is the number of patterns of the form P Op X that have yet to be expanded. This value is maintained for each operator Op and the pattern P. That is, a value for  $PNX$ ;  $PBKX$ ;  $XB_KP$ , etc. is maintained. X ranges over all events. For example, if there are 20 occurrences of  $P = \alpha B_K \beta$  in the input string and 5 patterns of the form  $\alpha B_K \beta NX$  have been discovered so far, then Actual Remaining Next  $\alpha B_K \beta NX = 15$ .

The following heuristic can be used to determine which discovered pattern is the best pattern to use for the expansion. Given an arbitrary literal D, the best pattern P for the expansion is preferably the pattern for which the value of

$E[[A[P \text{ Op } \delta]/E[P \text{ Op } \delta]]]$  is the maximum for some  $\delta$ .

This heuristic can be a probabilistic statement that the pattern P (which is most likely to result in the discovery of an interesting pattern) is the pattern for which there exists a literal  $\delta$ . In particular, the expected value of the interestingness measure of the pattern generated when the literal  $\delta$  is added to the pattern P via one of the temporal operators Op is the highest over all discovered patterns P, literals  $\delta$  and operators Op. It is preferable to use the expected value of the interestingness measure because,

although the actual number of occurrences of both the pattern  $P$  and the literals  $\delta$  is known, the number of occurrences of  $P \text{ Op } \delta$  is not known. This expectation is computed preferably directly from the previously-described derivations of expectations, and can be described using the following example:

5           If  $P = \alpha N \beta$ , and  $\text{Op}$  is "next",  
           then  $E[A[PN\delta]/E[PN\delta]]$   
                    $= (\#P_N)(FR(\delta)) = Pr[\alpha]Pr[\beta]Pr[\delta](K-2)$   
           where,  $K$  = length of input string,  
                    $FR(\delta)$  = frequency of the literals'  $\delta$  that could complete the  
 10           pattern  $\_N\_NX$ , and  
            $\#P_N$  = number of occurrences of the pattern  $P$  yet to be  
           expanded via the operator  $N$ .

If  $\text{Op}$  is "before", then

$E[A[PB_K\delta]/E[PB_K\delta]]$   
 15            $= ((\#P)(FR(\delta))(\text{"BEFOREK"}))/Pr[\alpha]Pr[\beta]Pr[\delta](K-2)(\text{"BEFOREK"})$   
            $= ((\#P)(FR(\delta)))/Pr[\alpha]Pr[\beta]Pr[\delta](K-2)$

where BEFOREK is a user defined variable that is equal to the  
 maximum distance between two events  $X$  and  $Y$  for  $XB_KY$  to hold.

Similar arguments can be used for any combination of the operators  $\text{Op}$  of "before",  
 20 "next", and "until". In general, the candidate pattern  $P$ , the suffix, the literal  $\delta$  and the  
 operator  $\text{Op}$  are chosen whose combinations are most likely to result in the discovery  
 of the interesting pattern.

Throughout the above-described technique and with reference to Figure 4B,  
 two data structures should be used to efficiently compute best candidates on each  
 25 subsequent iteration.

a. An  $((N+1) \times M)$  matrix where  $N$  is the number of distinct events, and  $M$  is the  
       number of different pattern forms that are intended to be discovered. For  
       example,  $M$  can be very large. However, it is preferable to limit the length of  
       the patterns to approximately 5 (depending on the application), taking into  
 30       consideration that the infrequency of much larger patterns typically makes  
       them statistically insignificant. With the maximum pattern length set to 5 and

-20-

using four temporal operators N, Bk, U, and ^, the value of

$$m = \sum_{i=1}^5 4^i = 4 \frac{(4^5 - 1)}{(4 - 1)} = 1364, \text{ which is a manageable number.}$$

The structure of this matrix can be follows: each entry  $[i, j]$   $i \in 1 \dots N, j \in 1 \dots M$  represents the remaining number of yet-to-be-discovered patterns

5 having the form  $j$  whose final event is  $i$ . This number can be easily maintained because it is the total number of occurrences of the event  $i$  minus the number of already discovered patterns of the form  $j$  whose final event is  $i$ . The additional  $(N+1)$  row contains the total number of already discovered patterns (i.e., the sum of the values in the columns) of the form  $j$ . Each column of this  
10 array can be sorted such that literal  $\alpha$  precedes  $\beta$  in the column  $j$  if the number of the literals  $\alpha$  remaining to be added as suffixes to create patterns of the form  $j$  divided by  $Pr[\alpha]$ , exceeds that value for the literal  $\alpha$ . This value can be called the "candidacy value" of the corresponding literal for the corresponding pattern form. The matrix can be called the "suffix matrix".

15 b. The second data structure is an array of  $M \times R$  lists where  $M$  is again number of different pattern forms that should be discovered and  $R$  is the number of temporal operators being used. In list  $j_{op}$ , all patterns of the form  $j$  that have already been discovered are maintained in a sorted order by the number of the occurrences of each pattern yet to be expanded through the use of the operator  
20 Op divided by  $E[P]$ . This value can be called the corresponding pattern's "candidacy value" for the corresponding operator. Such value is simple to calculate since the total number of patterns that have the form  $P \text{ Op } X$  is known. Along with each pattern, it is possible to maintain the number of occurrences of the given pattern  $P$ , and the locations of the pattern  $P$ . This  
25 array can be termed the "set of discovered patterns."

The best combination of an element from each of these two data structures may be the candidate for the next discovery iteration. For example, at each iteration, it is possible to assume that the first value in each list in the set of discovered patterns of whose length is less than the maximum allowed pattern length corresponds to the  
30 patterns  $P_1, P_2, \dots, P_M$ . Additionally, it is possible to assume that the first value in

each column in the suffix matrix may correspond to the literals  $\alpha_1, \alpha_2, \dots, \alpha_M$ . The  $M$  values that result from multiplying the candidacy value are computed for each of these patterns  $P_i$  times the first value in the suffix matrix for the pattern form that is the result of combining the pattern  $P_i$  from the set of discovered patterns with the  
5 literal  $\alpha$  via the operator  $Op$  corresponding to the operator for the list from which the pattern  $P$  was taken. The pattern  $P_i$ , literal  $\alpha_j$  and operator  $Op$  can be selected whose combination results in the largest value among these  $M$  values. In doing so, the goal of selecting the candidate pattern, literal, and operator whose combination is most likely to result in the discovery of an interesting pattern can be accomplished. Once  
10 these candidates have been selected, the determination of the number of occurrences of the pattern  $P_i Op \alpha_j$  can be computed via linear scans of the location lists for the pattern  $P_i$  and the literal  $\alpha_j$ . For example, if  $Op=N$ , then it is possible to look for locations  $l$  such that  $P_i$  occurs at the location  $l$  and  $\alpha_j$  occurs at location  $l+1$ . If  $Op = \wedge$ , it is possible to look for the locations where both  $P_i$  and  $\alpha_j$  occur. One of the ways to  
15 initiate the above-described procedure is by choosing the variable triple (i.e., pattern, literal, operator) whose combination would most likely result in the discovery of an interesting pattern. As the procedure progresses, if the given pattern  $P$  has not generated many newly discovered patterns as a candidate for the expansion, the pattern will preferably percolate toward the top of its associated sorted list. Likewise,  
20 if a literal  $\alpha$  has not been used as the suffix of many discovered patterns, the literal will percolate to the top of its suffix list. In this way, as patterns and literals become more likely to generate an interesting pattern, via the combination, and they will become more likely to be chosen as candidates for the next iteration.

## II. CLASSIFICATION OF DATA

25 Turning back to the method of the present invention illustrated in Figure 2, the data obtained in the feature selection step is classified (step 110). The classification of data has been problematic to those having ordinary skill in the art of data mining. The most widely utilized classification technique entails the use of decision trees. There are more powerful classification techniques (in the sense that the decision trees  
30 are able to represent a more robust class of functions) such as neural networks. However, those having ordinary skill in the art often do not use the neural networks

for classifying data because the neural networks are computationally complex, and lack transparency. One of the important features of a classifier is that the resulting function should ultimately be understandable. It is preferable to understand why a prediction made by the classifier was made to better understand relationships that exist in current problem domain. The neural networks are a black box, and while their predictions may be accurate, they lead to little insight about the problem at hand.

The present invention uses an alternative technique known as "MARS" (Multivariate Adaptive Regression Splines). The detailed description of MARS is described in, e.g., in J. Friedman, "Multivariate Adaptive Regression Splines", The Annals of Statistics, Vol. 19, No. 1, 1991 pp. 1-141. MARS is a nonlinear technique that overcomes many of the shortcomings of standard decision trees while being computationally tractable and ultimately interpretable.

Although the recursive partitioning may be the most adaptive of the methods for multivariate function approximation, it suffers from some significant restrictions that limit its effectiveness. One of the most significant of these restrictions is that the approximating function is discontinuous at the subregion boundaries (as defined by splits in the nodes). It severely limits the accuracy of the approximation, especially when the true underlying function is continuous. Small perturbations in the predictor variables can potentially result in widely varying predictions. Additionally, decision trees are poor at approximating several simple classes of functions such as linear and additive functions. The records obtained from the feature selection are augmented by a set of temporal features. For example, from the data records having 9 features to the data having 200 features (i.e., a high dimensional data).

Since all classification techniques generate models based on in sample data that are designed to perform well on out of sample data and because of the resultant high-dimensionality, the issue of over-fitting may occur as described below.

### **III. BIAS INCREASE VIA SHRINKAGE**

In all classification techniques, the introduction of additional degrees of freedom reduces the in sample error (bias) of the model while increasing the model variance. This frequently results in poor approximations of out of sample data. To address this problem, some classification methods include a technique for reducing

the model bias, typically via a reduction in the classification model's degrees of freedom. This reduction in degrees of freedom increases bias in the classification model, while reducing its variance and out-of-sample error.

The combination of the forecasts can be done by averaging resulting in a maximum likelihood estimator ("MLE"). To evaluate the applicability and usefulness of this approach, it is possible to consider the more general situation of trying to estimate a parameter  $\Theta$  by  $t(x)$ . For example, if  $E[t(x)] = \Theta$ , then  $t(x)$  can be an unbiased estimator of  $\Theta$  and a measure of the precision of this estimator may be  $E[t(x) - \Theta]^2$ , i.e., its variance. Instead, if  $E[t(x)] \neq \Theta$ , then  $t(x)$  is a biased estimator of  $\Theta$ . A measure of its precision is still  $E[t(x) - \Theta]^2$ , but because  $E[t(x)] \neq \Theta$ , this quantity is not the variance, and known as the mean squared error. Thus,

$$\begin{aligned} E([t(x) - \Theta]^2) &= E([t(x) - E[t(x)] + E[t(x)] - \Theta]^2) \\ &= E([t(x) - E[t(x)]]^2) + (E[t(x)] - \Theta)^2 + 2(E[t(x)] - \Theta) E[t(x) - E[t(x)]] \\ &= E([t(x) - E[t(x)]]^2) + (E[t(x)] - \Theta)^2 \\ &= \text{var}[t(x)] + (E[t(x)] - \Theta)^2 \\ &= \text{var}[t(x)] + [\text{Bias}(t)]^2 \end{aligned}$$

By sacrificing an increased bias for a decreased variance, it is possible to achieve a uniformly-smaller MLE. Stein's estimator, now known as Stein shrinkage, described in B. Efron et al., "Stein's Estimation Rule and its Competitors-An Empirical Bayes Approach", Journal of the American Statistical Assoc., Vol. 68, March-1973, pp. 117-130, was originally developed for the case of reducing bias in linear functions. The results of the Stein's estimator can be extended for the nonlinear case. For example, by "shrinking" the estimated parameters towards the sample mean, this approach mitigates the effects of non-stationarity by reducing the impact of deviations in the distributions of the estimator variables between in-sample and out-of-sample data.

Thereafter, in step 130 of Figure 2, the prediction model is generated from the data records on which feature selection was performed, and/or which were classified and then shrunk. Finally, in step 140, such prediction model and/or the classified and reduced data are output to the printer 30, display device 40, storage device 50 and/or communications network 60.



#### IV. EXEMPLARY APPLICABILITY OF THE PRESENT INVENTION

The system and method according to the present invention can be used in two exemplary settings, e.g., a network intrusion detection and a disease classification. Embodiments of the present invention for each of these exemplary settings are  
5 discussed below.

A. Network intrusion detection addresses the problem of detecting intrusions on a computer network. In summary, the training data may consist of a set of TCP/IP records that have been scored 0/1 depending on whether that connection was part of an attack as well as with the specific attack type. The intrusion detection system then  
10 learns features that distinguish normal from malicious network activity. These features then become the input to a classifier which when run on out-of-sample data scores each record based on the likelihood that it is part of an attack. Finally, the third stage is to combine the classifiers that result from training on many in sample training sets as well as to mitigate the problems of over-fitting and non-stationarity.

15 Figure 5 shows one such exemplary intrusion detection system ("IDS") 400 according to the present invention. First, data is collected in the form of log-files that consist of a sequence of records about activity on the network. The log files can be collected via a local area network 440 from an information server 410, an attached firewall 420, user workstations 430 and/or other sites. One record can be created for  
20 each connection that occurs. The information in each record may include time and date of the connection, the type of service provided, the source and destination ports, the source and destination IP addresses, the duration of the connection, and other data.

The IDS 400 described above serves two purposes, e.g., data collection and network activity monitoring, and intrusion identification. In serving these roles, the  
25 IDS 400 may include or be connected to a large database (e.g., the storage device 50) for data storage, and a computational engine for scoring individual network activity records based on their likelihood of being part of an intrusion. In the training phase and as illustrated in Figure 6, the IDS accumulates the data generated at the various monitoring points on the network (step 500). The aggregated data records are then  
30 scored manually, e.g., with a score of 1 indicating that the given record was part of a network attack and a score of 0 indicating that the record was part of normal activity.

This exemplary embodiment of the present invention may use, e.g., scored data generated by the Defense Department's Advance Research Project Agency (DARPA). Once collected, this data becomes the input to the IDS 400, as shown in further detail in Figure 7. The initial set of data records represents the input to the first stage of the technique, i.e., feature selection. In this stage a set of additional features (typically several hundred) are generated and added to each data record (step 510). The first set of data records are set as the current data records in step 520, and the current data records are input to the second stage of the technique - classification, i.e., MARS. MARS generates a functional model of the data capable of predicting intrusions on out-of-sample data based on the current data records (step 530). Then, in step 540, it is determined whether all sets of the modified data records were utilized. If not, in step 550, the next set of the modified data records is set as the current set of records, and the process returns to step 530 so that a number of functional models are generated. This set of models is then input into the final stage of the technique, i.e., shrinkage. Shrinkage results in the generation of a single model based on the aggregation of all of the predictor models generated (step 560). This is done in a way to mitigate the effects of non-stationarity in the data. This final model is then incorporated into the IDS 400. In the IDS 400, the model monitors network activity, identifying activity that is part of an (attempted) intrusion on the network. Concurrently, the IDS 400 may accumulate data records generated by the network monitors for use as future training data to the model. This allows the system and method of the present invention to continuously update itself based on changes in the types of activity occurring on the network.

B. In the disease classification, the main focus can be on cancer. Given that cancer results from changes in the DNA of healthy cells, the present invention provides an approach to cancer classification based on the gene expression. Both the cancer classification problem as well the class discovery problem are addressed by identifying discrepancies in gene expression between healthy and cancerous cells. It is then possible to evaluate the quality of the approach of the system and method according to the present invention to cancer classification by considering RNA samples from both healthy individuals as well as samples from patients from multiple

known cancer classes as identified by their histopathological appearance for accurately and consistently validating the diagnosis made by hematopathologists on the genetic grounds. This is achieved by training the system (as described below) on RNA samples that are properly labeled by their cancer class (or labeled as being healthy). By discovering the genetic differences among cancer classes, a predictive model of these classes is generated which can then be tested via cross validation and through testing on out of sample data, and a class discovery can be performed. For example, the system is trained on the same RNA samples. This time, however, these samples are unlabeled. The classes associated with each sample are discovered without a prior knowledge of this information. Additionally, novel classes within these samples are discovered.

As shown in Figure 8, healthy DNA and cancerous DNA can each be dyed different colors and hybridized on a micro-array containing thousands of genes expected to be relevant to cell growth and regulation. Through this process, the expression levels of these targeted genes can be compared between the healthy and cancerous cells. The cancer classifier then constructs a model capable of classifying future DNA samples as either healthy or cancerous. Additionally, DNA samples from two different cancer types can be hybridized and a model constructed that identifies the cancer type of an out-of-sample, cancerous DNA strand. Through this process, the system is first capable of determining whether or not a DNA sample is cancerous, and if it is then identifying the associated cancer type. These results improve the targeting of treatment to specific cancer types. Described below is a description of how to distinguish between healthy and cancerous DNA, although the process may not be identical for identifying specific cancer types.

The data collected from the micro-array is a set of gene expression levels for both normal and cancerous DNA in thousands of different genes. Once collected, this data becomes the input to the cancer classification system (CCS) (see diagram below). As shown in Figure 9, the set of expression levels represents the input to the first stage of the method and system according to the present invention, i.e., feature selection. In this stage a set of features (typically several hundred) are generated. These features represent relevant relationships between the expression levels of

-27-

different genes in terms of their ability to distinguish healthy from cancerous DNA. An example of a potential feature is, e.g.,  $\text{ExpressionLevel}(\text{Gene}\#32) > T$  AND  $\text{ExpressionLevel}(\text{Gene}\#656) > T$ . This feature provides that both the expression levels of gene number 32 and number 656 exceed some threshold, and may be included if it

5 represented a situation that is either highly correlated with healthy or highly correlated with cancerous DNA. Thus, such features are input to the second stage of the technique, i.e., MARS. MARS generates a functional model of the data capable of distinguishing between healthy and cancerous DNA on out-of-sample data. This process is typically executed several times on different training data sets, thus

10 generating several models. This set of models is then input into the final stage of the technique, i.e., shrinkage. Shrinkage results in the generation of a single model based on the aggregation of all of the predictor models generated. The combination of models is particularly relevant to cancer classification when attempting to build a model that differentiates between several cancer types. Models are initially

15 constructed to distinguish between pairs of cancer classes. Shrinkage then combines these models to create a single monolithic classifier capable of distinguishing between many different cancer classes.

One having ordinary skill in the art would clearly recognize that many other domains and applicable example in which data is temporal and/or non-stationary in

20 nature can benefit using this system and method for classification according to the present invention. Indeed, the present invention is in no way limited to the exemplary applications and embodiments thereof described above.

CLAIMS

1. A method for determining a feature of a particular pattern, comprising the steps of:
  - a) receiving data records;
  - 5 b) obtaining predetermined patterns that are associated with at least some of the data records;
  - c) extracting particular information from at least a subset of the received data records, the particular information being indicative of the particular pattern for at least some of the data records; and
  - 10 d) determining whether the particular pattern is an unexpected pattern based on the obtained predetermined patterns.
2. The method according to claim 1, wherein at least one record of the data records includes temporal data.
3. The method according to claim 1, wherein at least one record of the data  
15 records includes non-stationary data.
4. The method according to claim 1, wherein step (b) comprises the substeps of:
  - i. assigning a threshold, and
  - ii. correlating the data records into sets of patterns as a function of the threshold.
- 20 5. The method according to claim 4, wherein step (d) includes the substep of determining if the particular pattern corresponds to at least one pattern of the sets of patterns.
- 6 The method according to claim 5, wherein the unexpected pattern is established if the particular pattern does not correspond to any pattern of the sets of  
25 patterns.
7. The method according to claim 1, wherein the unexpected pattern is indicative of an interestingness measure in the predetermined pattern.

8. The method according to claim 1, wherein the data records include input sequences, and wherein step (d) comprises the substep of scanning the input sequences to determine an interestingness measure of at least one event in the input sequences.
- 5 9. The method according to claim 8, wherein step (d) comprises the substeps of:
- i. initializing a pattern list by inserting all events of the input sequences therein, and
  - ii. from all patterns in the pattern list, selecting a first pattern which has a largest interestingness measure.
- 10 10. The method according to claim 9, wherein the data records include a maximum allowable length value, and wherein step (d) comprises the substeps of:
- iii. expanding the first pattern to be a second pattern,
  - iv. if a length of the second pattern is greater than the maximum allowable value, adding the second pattern to the pattern list, and
  - 15 v. if a length of the second pattern is less than or equal to the maximum allowable value, subtracting the first pattern from the pattern list.
- 20 11. The method according to claim 10, wherein step (d) comprises the substep of repeating substeps (ii)-(v) until the pattern list becomes empty.
12. The method according to claim 11, further comprising the step of:
- e) outputting the particular pattern which includes the interestingness measure.
- 25 13. The method according to claim 8, wherein step (d) comprises the substeps of:
- i. initializing a pattern list by inserting all events of the input sequences therein,
  - ii. initializing at least one suffix list,

- iii. calculating locations of certain patterns of the input sequences,
  - iv. updating previously discovered patterns based on the calculated locations, and
  - v. updating the at least one suffix list using the certain patterns.
- 5 14. The method according to claim 13, wherein the data records include a maximum allowable length value, and wherein step (d) comprises the substep of:
- vi. if a length of the second pattern is greater than or equal to the maximum allowable value, repeating substeps (iii)-(v).
- 10 15. The method according to claim 1, wherein step (d) includes the substep of:
- i. generating further records by modifying the data records to include additional features.
16. The method according to claim 15, further comprising the step of:
- (f) generating a functional model using the further records.
- 15 17. The method according to claim 16, wherein substep (i) includes generating a plurality of sets of the further records, and wherein step (f) is executed for each set of the further records.
18. The method according to claim 17, wherein step (f) includes the substep of generating a single model based on each functional model of the respective set of the
- 20 further records.
19. The method according to claim 1, further comprising the steps of:
- (g) after step (d), classifying the data records which have the unexpected pattern associated therewith; and
  - (h) generating a prediction model as a function of the classified data
- 25 records.
20. The method according to claim 19, wherein step (g) is performed using a Multivariate Adaptive Regression Splines technique.

21. The method according to claim 19, further comprising the step of:  
(i) shrinking at least one of data and parameters of the classified data records.
22. The method according to claim 21, wherein step (i) includes the substep of  
5 determining a mean of the at least one of the data and the parameters.
23. The method according to claim 21, wherein step (i) is performed using a Stein's Estimator Rule technique.
24. The method according to claim 1, wherein at least one of the predetermined patterns utilizes temporal modal operators.
- 10 25. The method according to claim 1, wherein at least one of the predetermined patterns utilizes logical connectives.
26. The method according to claim 1, wherein at least one of the predetermined patterns is generated by a computer program.
27. A system for determining a feature of a particular pattern, comprising:  
15 a processing arrangement programmed to:  
a) receiving data records,  
b) obtaining predetermined patterns that are associated with at least some  
of the data records,  
c) extracting particular information from at least a subset of the received  
20 data records, the particular information being indicative of the  
particular pattern for at least some of the data records, and  
d) determining whether the particular pattern is an unexpected pattern  
based on the obtained predetermined patterns.
28. The system according to claim 27, wherein at least one record of the data  
25 records includes temporal data.
29. The system according to claim 27, wherein at least one record of the data records includes non-stationary data.



30. The system according to claim 27, wherein, in step (b), the processing arrangement:

- i. assigns a threshold, and
- ii. correlates the data records into sets of patterns as a function of the threshold.

5

31. The system according to claim 30, wherein, in step (d), the processing arrangement determines if the particular pattern corresponds to at least one pattern of the sets of patterns.

32. The system according to claim 31, wherein the unexpected pattern is established if the particular pattern does not correspond to any pattern of the sets of patterns.

10

33. The system according to claim 27, wherein the unexpected pattern is indicative of an interestingness measure in the predetermined pattern.

~~34. The system according to claim 27, wherein the data records include input sequences, and wherein step (d) comprises the substep of scanning the input sequences to determine an interestingness measure of at least one event in the input sequences.~~

15

35. The system according to claim 34, wherein, in step (d), the processing arrangement:

20

- i. initializes a pattern list by inserting all events of the input sequences therein, and
- ii. from all patterns in the pattern list, selects a first pattern which has a largest interestingness measure.

36. The system according to claim 35, wherein the data records include a maximum allowable length value, and wherein, in step (d), the processing arrangement:

25

- iii. expands the first pattern to be a second pattern,

-33-

- iv. if a length of the second pattern is greater than the maximum allowable value, adds the second pattern to the pattern list, and
- v. if a length of the second pattern is less than or equal to the maximum allowable value, subtracts the first pattern from the pattern list.
- 5

37. The system according to claim 36, wherein, in step (d), the processing arrangement repeats substeps (ii)-(v) until the pattern list becomes empty.

38. The system according to claim 37, wherein the processing arrangement is further programmed to:

- 10 e) output the particular pattern which includes the interestingness measure.

39. The system according to claim 34, wherein, in step (d), the processing arrangement:

- i. initializes a pattern list by inserting all events of the input sequences therein,
- 15 ii. initializes at least one suffix list,
- iii. calculates locations of certain patterns of the input sequences,
- iv. updates previously discovered patterns based on the calculated locations, and
- 20 v. updates the at least one suffix list using the certain patterns.

40. The system according to claim 39, wherein the data records include a maximum allowable length value, and wherein, in step (d), the processing arrangement:

- vi. repeats substeps (iii)-(v) if a length of the second pattern is greater than or equal to the maximum allowable value.
- 25

41. The system according to claim 27, wherein, in step (d), the processing arrangement:

-34-

- i. generates further records by modifying the data records to include additional features.

42. The system according to claim 41, wherein the processing arrangement is further programmed to:

- 5 (f) generates a functional model using the further records.

43. The system according to claim 42, wherein, in substep (i), the processing arrangement generates a plurality of sets of the further records, and wherein the processing arrangement executes step (f) for each set of the further records.

44. The system according to claim 43, wherein, in step (f), the processing  
10 arrangement generates a single model based on each functional model of the respective set of the further records.

45. The system according to claim 27, wherein the processing arrangement is further programmed to:

- (g) after step (d), classify the data records which have the unexpected  
15 pattern associated therewith, and

- (h) generate a prediction model as a function of the classified data records.

46. The system according to claim 45, wherein the processing arrangement performs step (g) using a Multivariate Adaptive Regression Splines technique.

47. The system according to claim 45, wherein the processing arrangement is  
20 further programmed to:

- (i) shrink at least one of data and parameters of the classified data records.

48. The system according to claim 47, wherein, in step (i), the processing arrangement determines a mean of the at least one of the data and the parameters.

49. The system according to claim 47, wherein the processing arrangement  
25 performs step (i) using a Stein's Estimator Rule technique.

-35-

50. The system according to claim 27, wherein at least one of the predetermined patterns utilizes temporal modal operators.

51. The system according to claim 27, wherein at least one of the predetermined patterns utilizes logical connectives.

5 52. The system according to claim 27, wherein at least one of the predetermined patterns is generated by a computer program.

53. A method for classifying and reducing at least one of data and parameters provided in the data records, comprising the steps of:

- a) receiving data records;
- 10 b) classifying the data records which have at least one particular pattern, the data records being classified using a Multivariate Adaptive Regression Splines technique; and
- c) shrinking the at least one of the data and the parameters of the classified data records using a Stein's Estimator Rule technique.

15 54. The method according to claim 53, further comprising the steps of:

- d) obtaining predetermined patterns that are associated with at least some of the data records;
- e) extracting particular information from at least a subset of the received data records, the particular information being indicative of the at least one particular
- 20 pattern in at least some of the data records; and
- f) determining whether the at least one particular pattern is an unexpected pattern based on the obtained predetermined patterns.

55. A system for classifying and reducing at least one of data and parameters provided in data records, comprising:

- 25 a processing arrangement programmed to:
  - a) receive the data records,

-36-

- b) classify the data records which have at least one particular pattern, the data records being classified using a Multivariate Adaptive Regression Splines technique, and
- c) shrink at least one of data and parameters of the classified data records using a Stein's Estimator Rule technique.

5

56. The system according to claim 55, wherein the processing arrangement is further programmed to:

- d) obtain predetermined patterns that are associated with at least some of the data records,
- e) extract particular information from at least a subset of the received data records, the particular information being indicative of the at least one particular pattern in at least some of the data records, and
- f) determine whether the at least one particular pattern is an unexpected pattern based on the obtained predetermined patterns.

10

---

1/10




FIG. 1

2/10

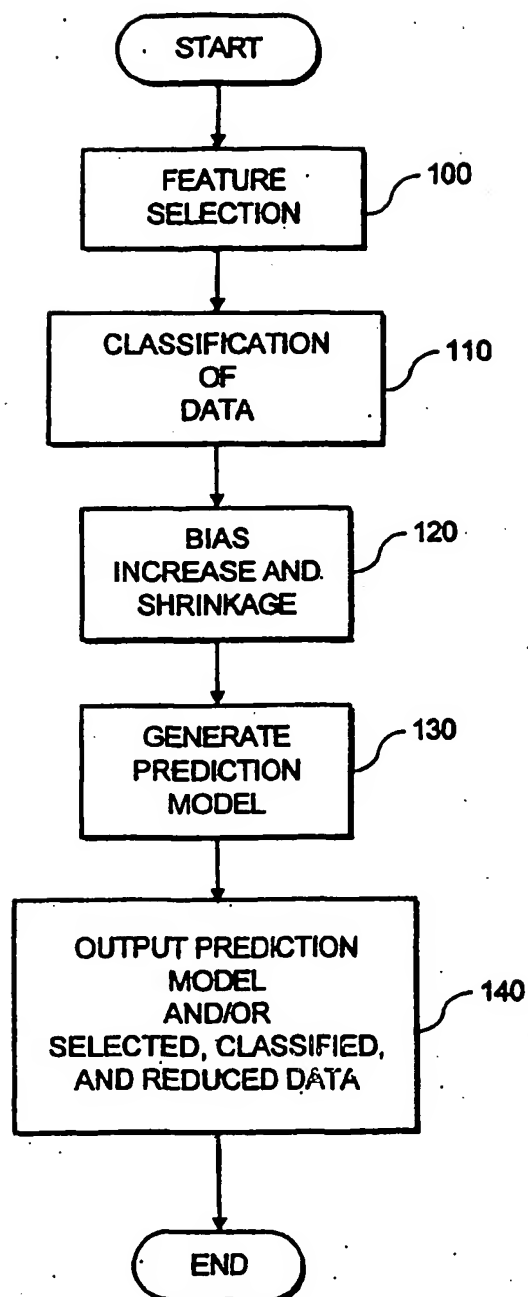


FIG. 2



WO0188834

Biblio

Desc

Claims

Page 1

Drawing

Prev

Next

*espace*



3/10

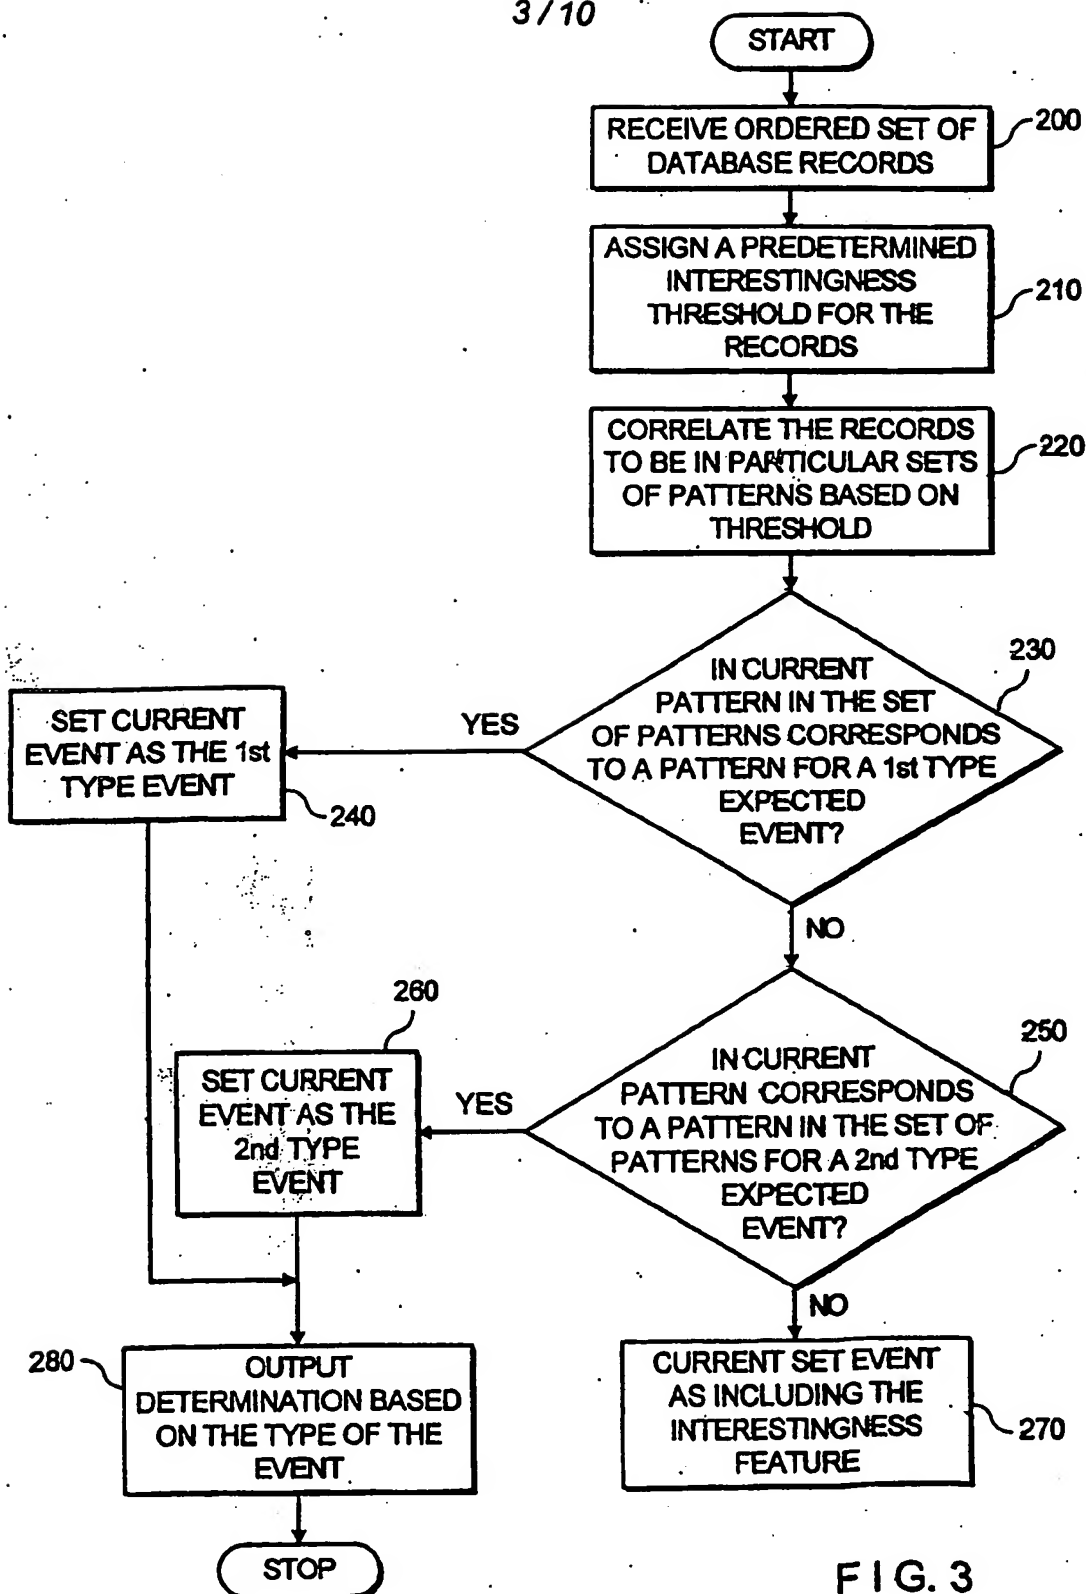


FIG. 3



WO0188834

Biblio

Desc

Claims


Page 1

Drawing

Prev

Next

*espace*





WO0188834

Biblio

Desc

Claims


Page 1

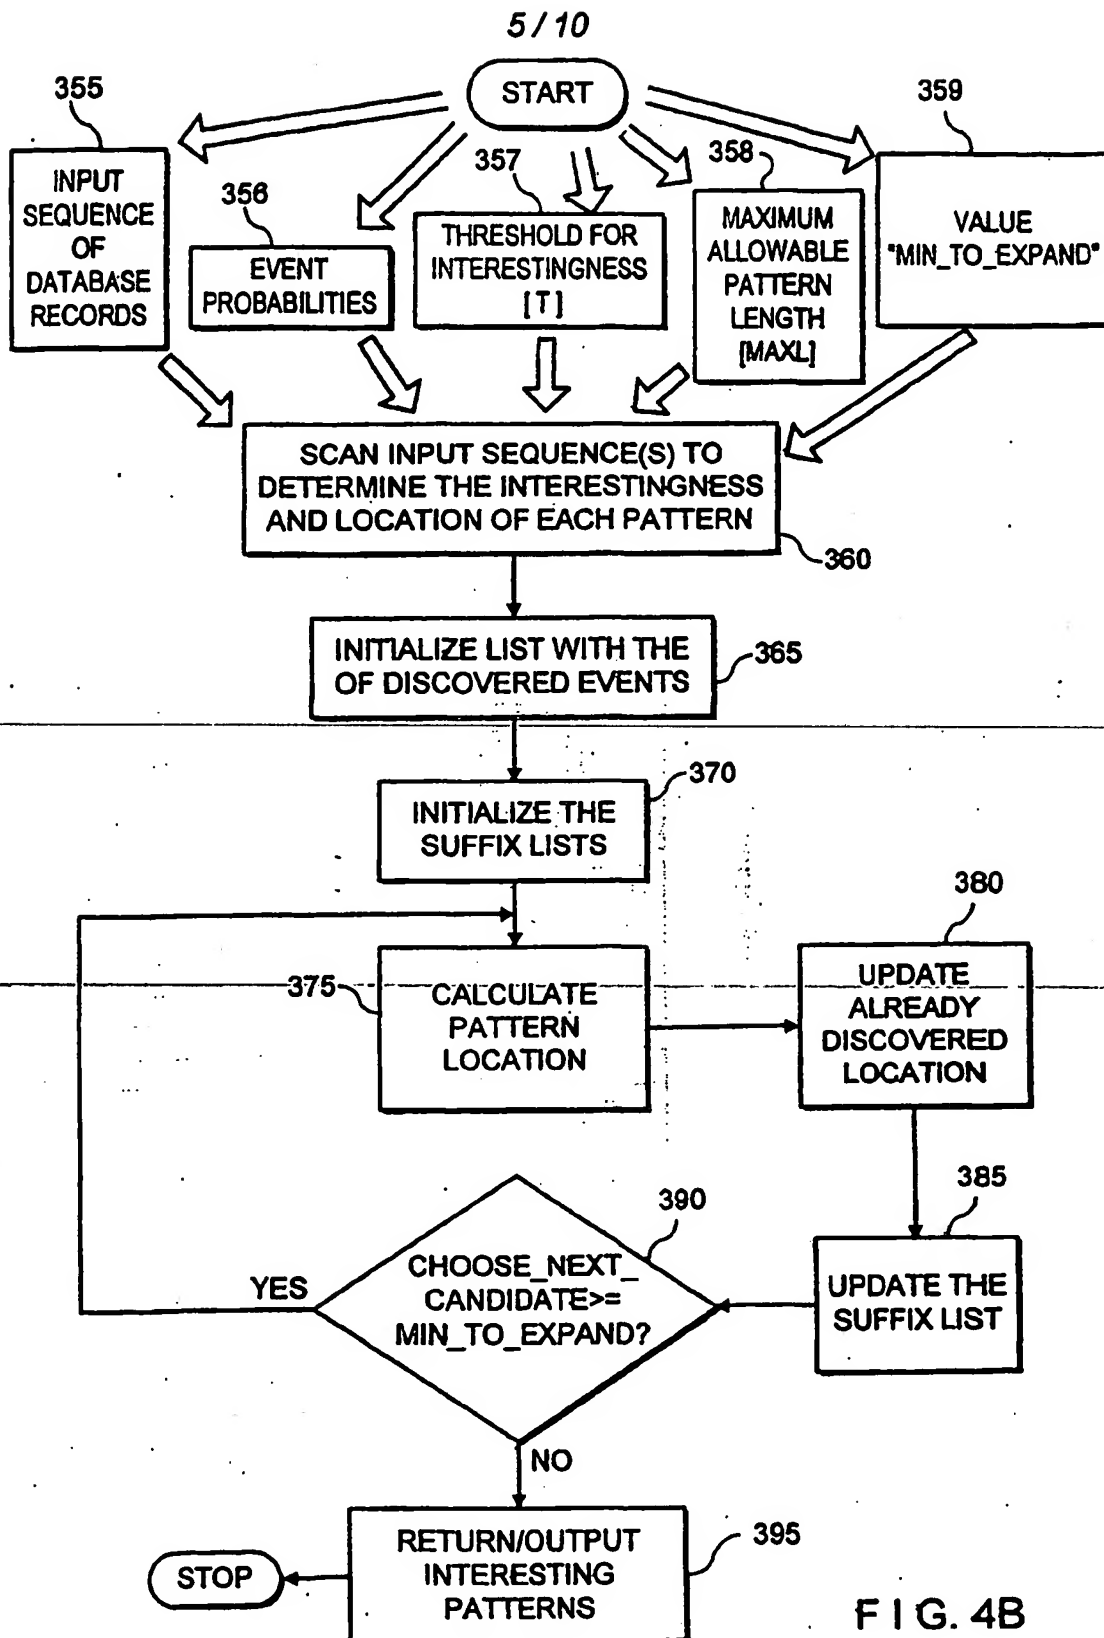
Drawing

Prev

Next

*espaçen*





6 / 10


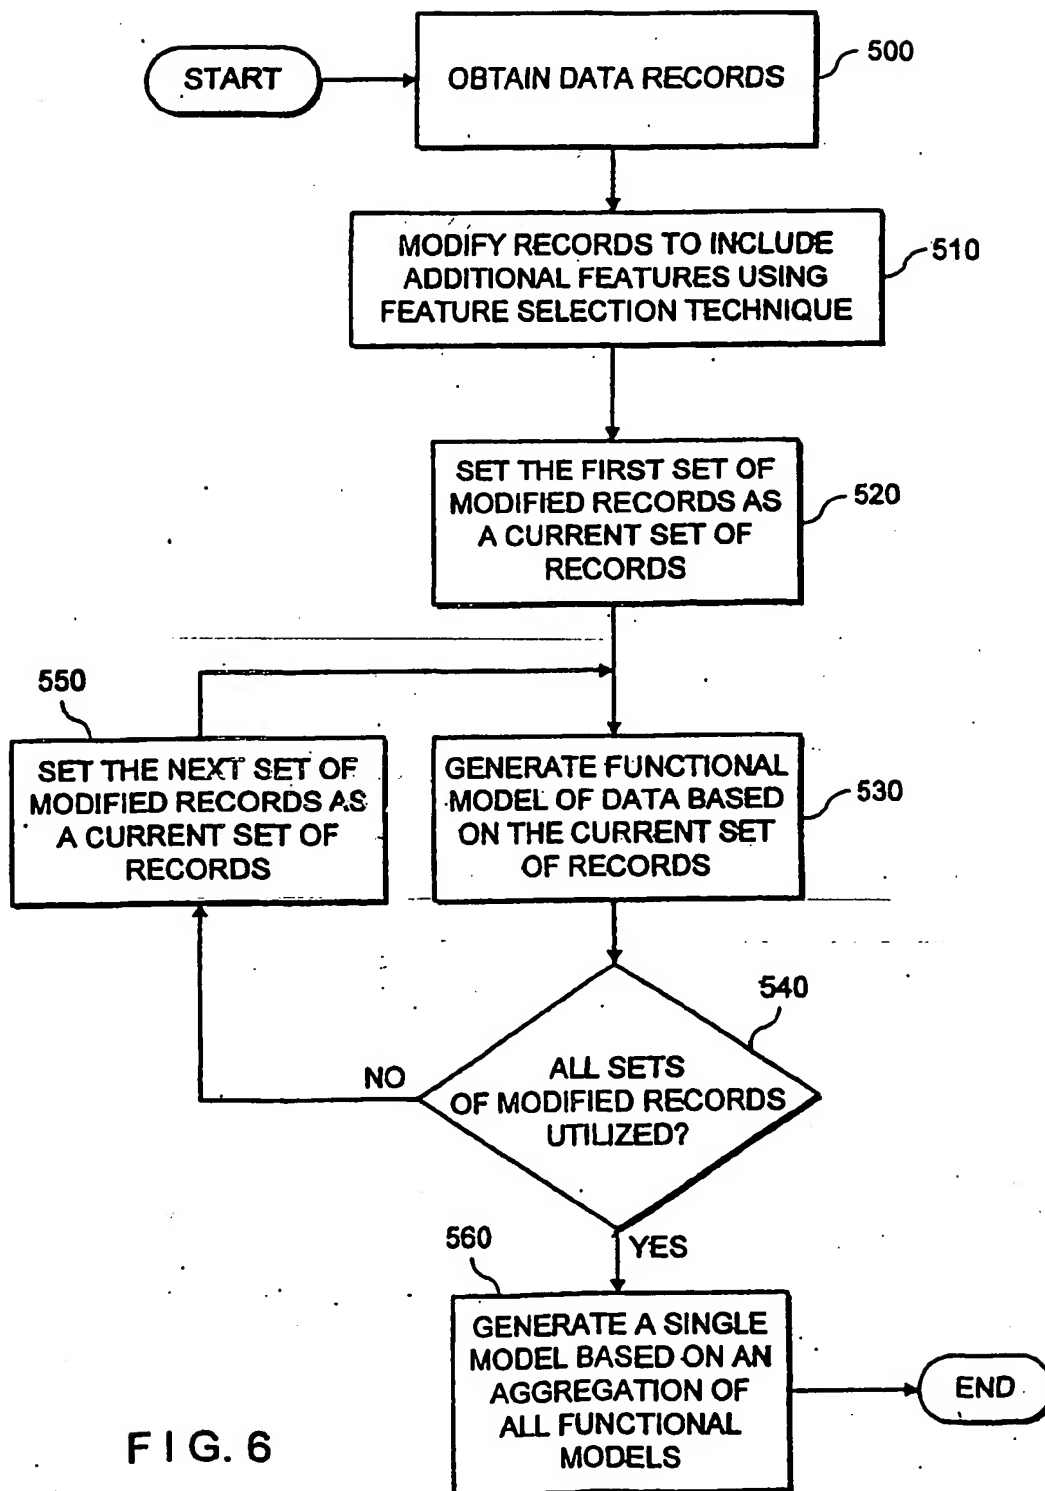


FIG. 5

7/10



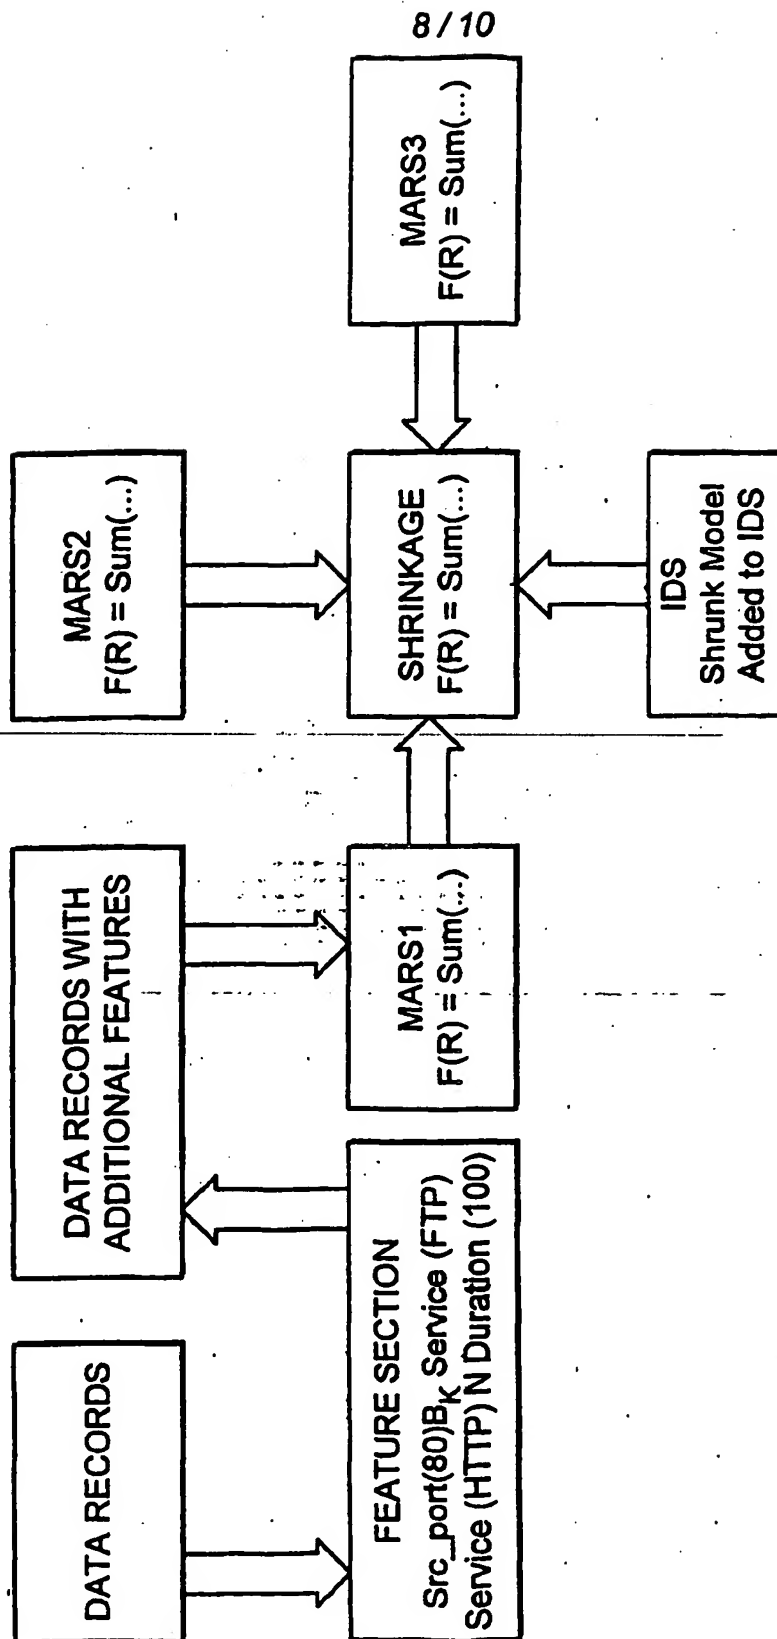


FIG. 7

9/10

## Gene copy number chips




FIG. 8



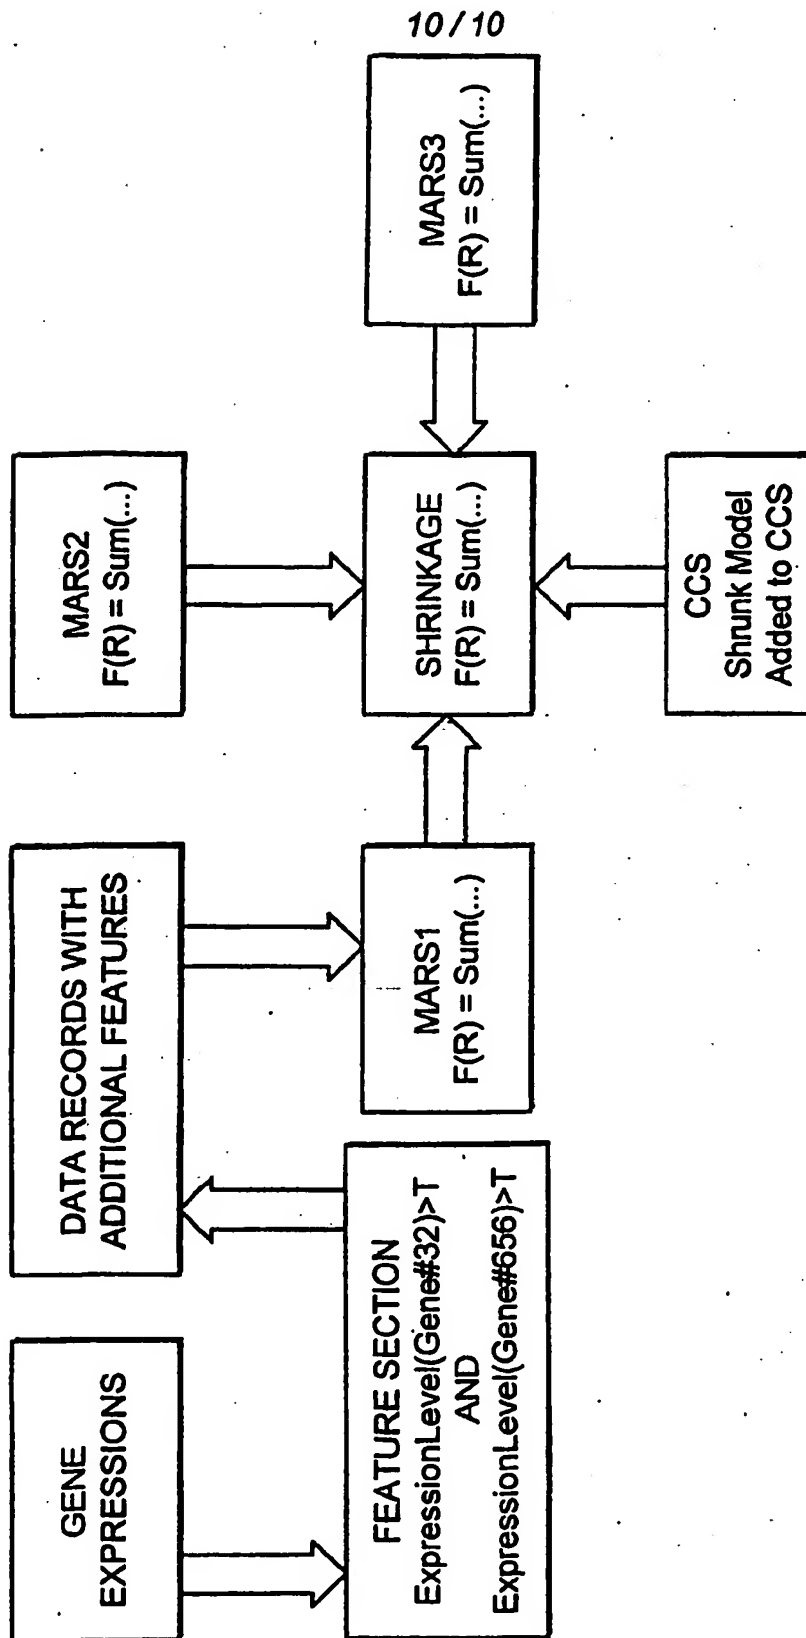


FIG. 9